



US009442941B1

(12) **United States Patent**
Luz et al.

(10) **Patent No.:** **US 9,442,941 B1**
(45) **Date of Patent:** **Sep. 13, 2016**

- (54) **DATA STRUCTURE FOR HASH DIGEST METADATA COMPONENT**
- (71) Applicant: **EMC Corporation**, Hopkinton, MA (US)
- (72) Inventors: **Kobi Luz**, Rehovot (IL); **Tal Ben-Moshe**, Kiryat Ono (IL); **Renan Hallak**, Sde Warburg (IL)
- (73) Assignee: **EMC CORPORATION**, Hopkinton, MA (US)
- (*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 178 days.

2005/0144416	A1	6/2005	Lin	
2005/0193084	A1	9/2005	Todd et al.	
2006/0031653	A1	2/2006	Todd et al.	
2009/0222596	A1	9/2009	Flynn et al.	
2010/0042790	A1	2/2010	Mondal et al.	
2010/0180145	A1	7/2010	Chu	
2012/0124282	A1*	5/2012	Frank	G06F 3/061 711/108
2013/0290285	A1*	10/2013	Gopal	G06F 17/3033 707/698
2014/0019764	A1*	1/2014	Gopal	H04L 9/3247 713/176
2014/0032992	A1*	1/2014	Hara	G06F 17/30156 714/773
2014/0244598	A1*	8/2014	Haustein	G06F 17/30156 707/692
2015/0019507	A1*	1/2015	Aronovich	G06F 17/30156 707/692
2015/0098563	A1*	4/2015	Gulley	H04L 9/0643 380/28

(21) Appl. No.: **14/229,491**

FOREIGN PATENT DOCUMENTS

(22) Filed: **Mar. 28, 2014**

EP	1804157	7/2007
WO	WO 2010/019596	2/2010
WO	WO 2010/040078	4/2010
WO	WO 2012/066528	5/2012

(51) **Int. Cl.**
G06F 17/00 (2006.01)
G06F 17/30 (2006.01)

OTHER PUBLICATIONS

(52) **U.S. Cl.**
CPC **G06F 17/30097** (2013.01)

U.S. Appl. No. 12/945,915, downloaded Mar. 27, 2014.
(Continued)

(58) **Field of Classification Search**
None
See application file for complete search history.

Primary Examiner — Hung Le

(74) *Attorney, Agent, or Firm* — Daly, Crowley, Mofford & Durkee, LLP

(56) **References Cited**


U.S. PATENT DOCUMENTS

5,860,137	A	1/1999	Raz et al.	
6,085,198	A *	7/2000	Skinner	G06F 8/315
6,125,399	A *	9/2000	Hamilton	H04L 29/06 370/392
6,643,654	B1	11/2003	Patel et al.	
7,908,436	B1	3/2011	Srinivasan et al.	
7,908,484	B2 *	3/2011	Haukka	H04L 9/0844 380/28
8,200,923	B1	6/2012	Healey et al.	
8,478,951	B1	7/2013	Healey et al.	
9,104,326	B2 *	8/2015	Frank	G06F 3/061
2005/0125626	A1	6/2005	Todd	

(57) **ABSTRACT**

In one aspect, a method includes providing a server that stores a hash digest metadata table and configuring entries in the hash digest metadata table to include entries that include persistent fields and non-persistent fields. The hash digest metadata table maps a hash digest representing X-Page data to its metadata information. The persistent fields include a short hash digest field, a physical layout offset field, a deduplication reference count field and a flags field.

14 Claims, 9 Drawing Sheets



(56)

References Cited

OTHER PUBLICATIONS

Advance E-Mail Notification Concerning Transmittal of International Preliminary Report on Patentability dated May 30, 2013 from the ISA for International Appl. No. PCT/IL2011/000692.

International Search Report and the Written Opinion of the ISA dated Dec. 1, 2011 for PCT Appl. No. PCT/IL2011/000692.

U.S. Appl. No. 13/901,062, filed May 23, 2013, entitled "Method & Apparatus for Block Level Data De-Duplication".

* cited by examiner




FIG. 1




FIG. 2

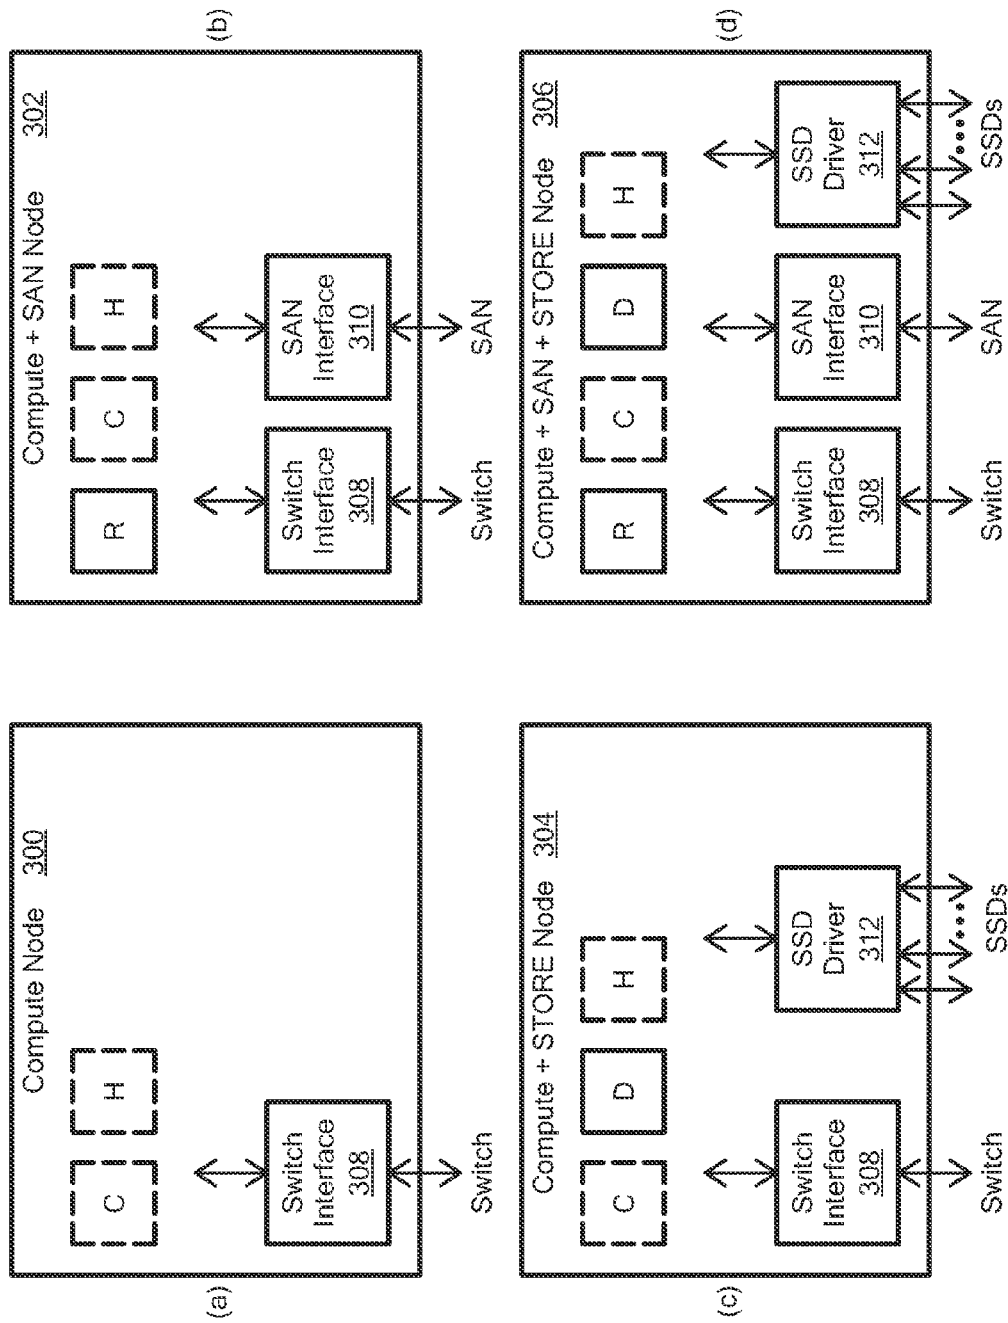




FIG. 3

**FIG. 4****FIG. 5**




FIG. 6

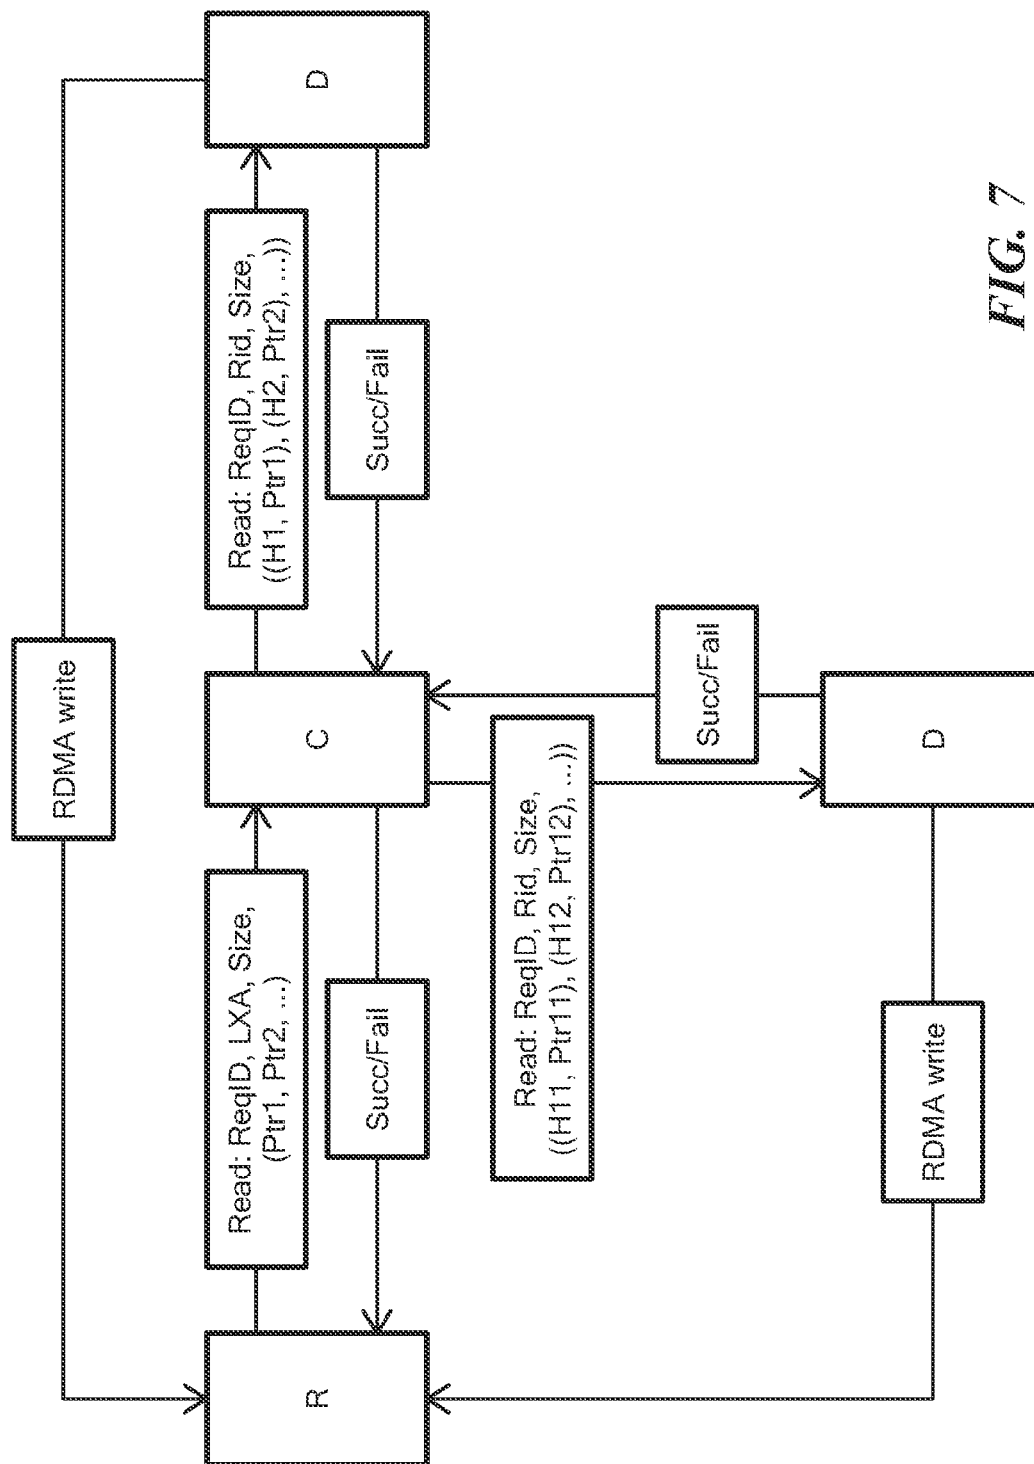


FIG. 7

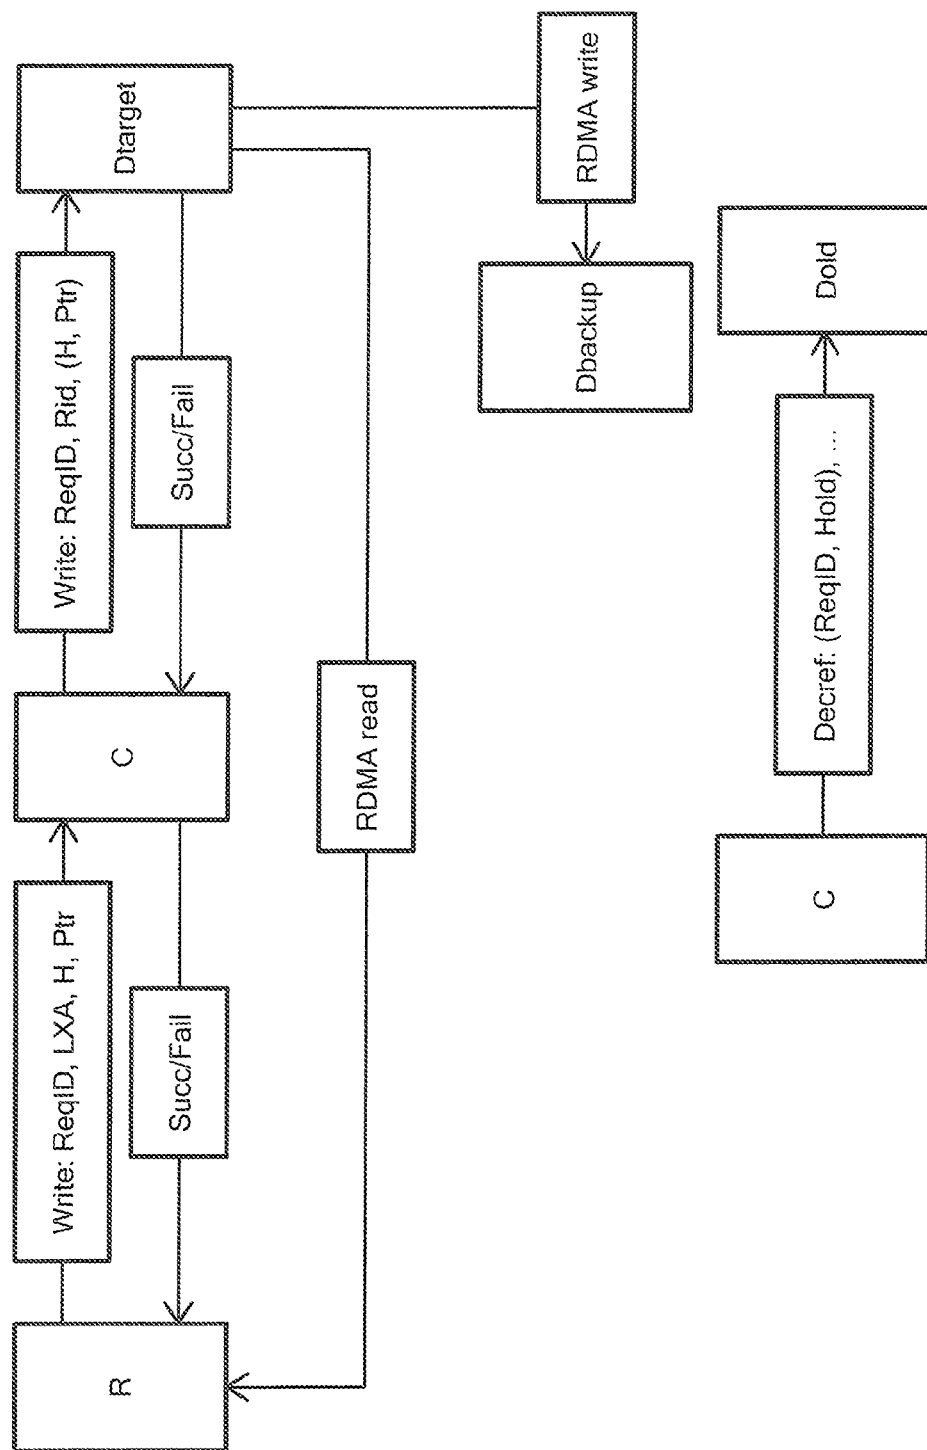


FIG. 8

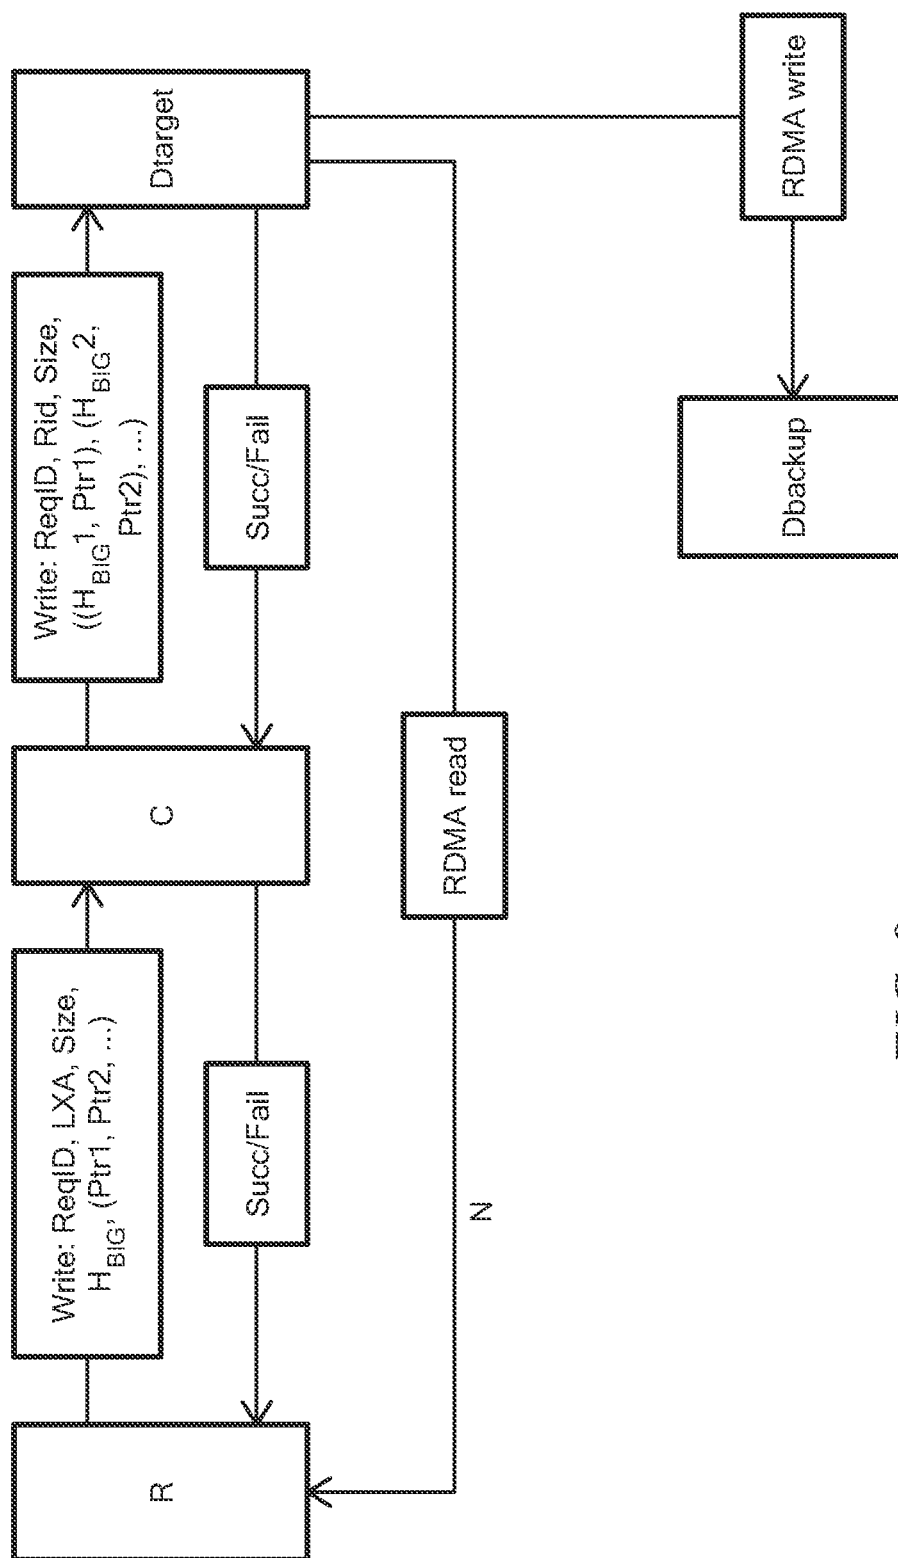


FIG. 9




FIG. 10




FIG. 11

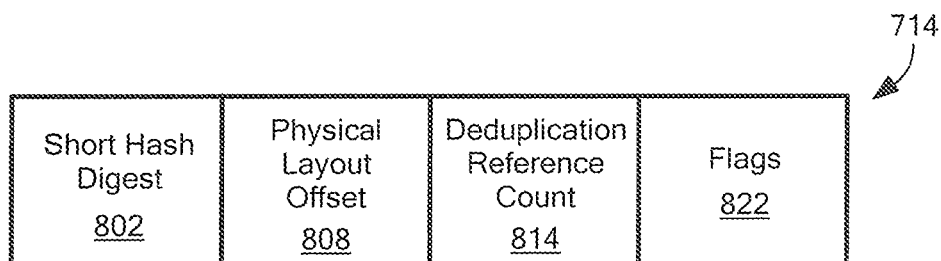


FIG. 12

DATA STRUCTURE FOR HASH DIGEST METADATA COMPONENT

BACKGROUND

Storage systems in general, and block based storage systems specifically, are a key element in modern data centers and computing infrastructure. These systems are designed to store and retrieve large amounts of data, by providing data block address and data block content—for storing a block of data—and by providing a data block address for retrieval of the data block content that is stored at the specified address.

Storage solutions are typically partitioned into categories based on a use case and application within a computing infrastructure, and a key distinction exists between primary storage solutions and archiving storage solutions. Primary storage is typically used as the main storage pool for computing applications during application run-time. As such, the performance of primary storage systems is very often a key challenge and a major potential bottleneck in overall application performance, since storage and retrieval of data consumes time and delays the completion of application processing. Storage systems designed for archiving applications are much less sensitive to performance constraints, as they are not part of the run-time application processing.

In general computer systems grow over their lifetime and the data under management tends to grow over the system lifetime. Growth can be exponential, and in both primary and archiving storage systems, exponential capacity growth typical in modern computing environment presents a major challenge as it results in increased cost, space, and power consumption of the storage systems required to support ever increasing amounts of information.

Existing storage solutions, and especially primary storage solutions, rely on address-based mapping of data, as well as address-based functionality of the storage system's internal algorithms. This is only natural since the computing applications always rely on address-based mapping and identification of data they store and retrieve. However, a completely different scheme in which data, internally within the storage system, is mapped and managed based on its content instead of its address has many substantial advantages. For example, it improves storage capacity efficiency since any duplicate block data will only occupy actual capacity of a single instance of that block. As another example, it improves performance since duplicate block writes do not need to be executed internally in the storage system. Existing storage systems, either primary storage systems or archiving storage systems are incapable of supporting the combination of content based storage—with its numerous advantages—and ultra-high performance. This is a result of the fact that the implementation of content based storage scheme faces several challenges:

(a) intensive computational load which is not easily distributable or breakable into smaller tasks,

(b) an inherent need to break large blocks into smaller block sizes in order to achieve content addressing at fine granularity. This block fragmentation dramatically degrades the performance of existing storage solutions,

(c) inability to maintain sequential location of data blocks within the storage systems, since mapping is not address based any more, and such inability causes dramatic performance degradation with traditional spinning disk systems,

(d) the algorithmic and architectural difficulty in distributing the tasks associated with content based mapping over

a large number of processing and storage elements while maintaining single content-addressing space over the full capacity range of the storage system.

A number of issues arise with respect to such devices, and it is necessary to consider such issues as performance, lifetime and resilience to failure of individual devices, overall speed of response and the like.

Such devices may be used in highly demanding circumstances where failure to process data correctly can be extremely serious, or where large scales are involved, and where the system has to be able to cope with sudden surges in demand.

SUMMARY

In one aspect, a method includes providing a server that stores a hash digest metadata table and configuring entries in the hash digest metadata table to include entries that include persistent fields and non-persistent fields. The hash digest metadata table maps a hash digest representing X-Page data to its metadata information. The persistent fields include a short hash digest field, a physical layout offset field, a deduplication reference count field and a flags field.

In another aspect, an article includes a non-transitory computer-readable medium that stores computer-executable instructions. The instructions cause a machine to provide a server that stores a hash digest metadata table. The hash digest metadata table maps a hash digest representing X-Page data to its metadata information. The hash digest metadata table includes entries that include persistent fields and non-persistent fields. The persistent fields include a short hash digest field, a physical layout offset field, a deduplication reference count field and a flags field.

In a further aspect, an apparatus includes electronic hardware circuitry configured to provide a server that stores a hash digest metadata table. The hash digest metadata table maps a hash digest representing X-Page data to its metadata information. The hash digest metadata table includes entries that include persistent fields and non-persistent fields. The persistent fields include a short hash digest field, a physical layout offset field, a deduplication reference count field and a flags field.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a simplified diagram schematically illustrating a system for data storage, having separate control and data planes.

FIG. 2 shows an exemplary configuration of modules for the system of FIG. 1.

FIG. 3 is a simplified diagram schematically illustrating four different node configurations for the system.

FIG. 4 is a simplified schematic diagram showing the nodes of FIG. 3 connected to a switch.

FIG. 5 is a simplified diagram showing a compute+SAN+store node for the device of FIG. 1.

FIG. 6 is a simplified flow chart which illustrates the flow for a read operation for one aligned X-page.

FIG. 7 is a simplified flow chart which illustrates the flow in the event that a read request arrives for a range of addresses spanning more than one X-Page but only one SL.

FIG. 8 is a simplified flow diagram illustrating the write procedure for a single aligned X page.

FIG. 9 is a simplified flow diagram illustrating the process for performing write operations to multiple full X-Pages.

FIG. 10 is a block diagram of an example of a data (D) module.

FIG. 11 is a block diagram of an example of a metadata entry including persistent fields and non-persistent metadata fields.

FIG. 12 is a block diagram of an example of the persistent fields in the metadata entry of FIG. 11.

DETAILED DESCRIPTION

A hash digest metadata table maps each in use hash digest, that represents actual data from an X-Page (as defined herein), to its metadata information including its physical page on the storage media (SSD), its memory copy (if it exists), a mapping to any backup memory copy and a reference count for the purpose of deduplication. Described herein is an approach to reduce the amount of memory consumed per hash digest and reach a point where most metadata is held in memory and requiring minimal reading of metadata from SSD as part of an I/O flow.

In a Content Addressable Storage (CAS) array, data is stored in blocks, for example of 4 KB, where each block has a unique large hash signature, for example of 20 bytes, saved on Flash memory.

The examples described herein include a networked memory system. The networked memory system includes multiple memory storage units arranged for content addressable storage of data. The data is transferred to and from the storage units using separate data and control planes. Hashing is used for the content addressing, and the hashing produces evenly distributed results over the allowed input range. The hashing defines the physical addresses so that data storage makes even use of the system resources.

A relatively small granularity may be used, for example with a page size of 4 KB, although smaller or larger block sizes may be selected at the discretion of the skilled person. This enables the device to detach the incoming user access pattern from the internal access pattern. That is to say the incoming user access pattern may be larger than the 4 KB or other system-determined page size and may thus be converted to a plurality of write operations within the system, each one separately hashed and separately stored.

Content addressable data storage can be used to ensure that data appearing twice is stored at the same location. Hence unnecessary duplicate write operations can be identified and avoided. Such a feature may be included in the present system as data deduplication. As well as making the system more efficient overall, it also increases the lifetime of those storage units that are limited by the number of write/erase operations.

The separation of Control and Data may enable a substantially unlimited level of scalability, since control operations can be split over any number of processing elements, and data operations can be split over any number of data storage elements. This allows scalability in both capacity and performance, and may thus permit an operation to be effectively balanced between the different modules and nodes.

The separation may also help to speed the operation of the system. That is to say it may speed up Writes and Reads. Such may be due to:

(a) Parallel operation of certain Control and Data actions over multiple Nodes/Modules

(b) Use of optimal internal communication/networking technologies per the type of operation (Control or Data), designed to minimize the latency (delay) and maximize the throughput of each type of operation.

Also, separation of control and data paths may allow each Control or Data information unit to travel within the system

between Nodes or Modules in the optimal way, meaning only to where it is needed and if/when it is needed. The set of optimal where and when coordinates is not the same for control and data units, and hence the separation of paths ensures the optimization of such data and control movements, in a way which is not otherwise possible. The separation is important in keeping the workloads and internal communications at the minimum necessary, and may translate into increased optimization of performance.

De-duplication of data, meaning ensuring that the same data is not stored twice in different places, is an inherent effect of using Content-Based mapping of data to D-Modules and within D-Modules.

Scalability is inherent to the architecture. Nothing in the architecture limits the number of the different R, C, D, and H modules which are described further herein. Hence any number of such modules can be assembled. The more modules added, the higher the performance of the system becomes and the larger the capacity it can handle. Hence scalability of performance and capacity is achieved.

The principles and operation of an apparatus and method according to the present invention may be better understood with reference to the drawings and accompanying description.

Reference is now made to FIG. 1 which illustrates a system 10 for scalable block data storage and retrieval using content addressing. The system 10 includes data storage devices 12 on which the data blocks are stored. The storage devices 12 are networked to computing modules, there being several kinds of modules, including control modules 14 and data modules 16. The modules carry out content addressing for storage and retrieval, and the network defines separate paths or planes, control paths or a control plane which goes via the control modules 14 and data paths or a data plane which goes via the data modules 16.

The control modules 14 may control execution of read and write commands. The data modules 16 are connected to the storage devices and, under control of a respective control module, pass data to or from the storage devices. Both the C and D modules may retain extracts of the data stored in the storage device, and the extracts may be used for the content addressing. Typically the extracts may be computed by cryptographic hashing of the data, as will be discussed in greater detail below, and hash modules (FIG. 2) may specifically be provided for this purpose. That is to say the hash modules calculate hash values for data which is the subject of storage commands, and the hash values calculated may later be used for retrieval.

Routing modules 18 may terminate storage and retrieval operations and distribute command parts of any operations to control modules that are explicitly selected for the operation in such a way as to retain balanced usage within the system 10.

The routing modules may use hash values, calculated from data associated with the operations, to select the control module for the distribution. More particularly, selection of the control module may use hash values, but typically relies on the user address and not on the content (hash). The hash value is, however, typically used for selecting the Data (D) module, and for setting the physical location for data storage within a D module.

The storage devices may be solid state random access storage devices, as opposed to spinning disk devices; however disk devices may be used instead or in addition.

A deduplication feature may be provided. The routing modules and/or data modules may compare the extracts or hash values of write data with hash values of already stored

data, and where a match is found, simply point to the matched data and avoid rewriting.

The modules are combined into nodes 20 on the network, and the nodes are connected over the network by a switch 22.

The use of content addressing with multiple data modules selected on the basis of the content hashing, and a finely-grained mapping of user addresses to Control Modules allow for a scalable distributed architecture.

A glossary is now given of terms used in the following description:

X-PAGE—A predetermined-size aligned chunk as the base unit for memory and disk operations. Throughout the present description the X-Page size is referred to as having 4 KB, however other smaller or larger values can be used as well and nothing in the design is limited to a specific value.

LUN or LOGICAL UNIT NUMBER, is a common name in the industry for designating a volume of data, or a group of data blocks being named with the LUN. Each data block is referred to, by the external user of the storage system, according to its LUN, and its address within this LUN

LOGICAL X-PAGE ADDRESS—Logical address of an X-Page. The address contains a LUN identifier as well as the offset of the X-Page within the LUN.

LOGICAL BLOCK—512 bytes (sector) aligned chunk, which is the SCSI base unit for disk operations.

LOGICAL BLOCK ADDRESS—Logical address of a Logical Block. The logical block address contains a LUN identifier as well as the offset of the logical block within the LUN.

SUB-LUN—Division of a LUN to smaller logical areas, to balance the load between C modules. Each such small logical area is called a sub-LUN.

SUB-LUN UNIT SIZE—The fixed size of a sub-LUN. **X-Page Data**—Specific sequence of user data values that resides in an X-Page. Each such X-Page Data is uniquely represented in the system by its hash digest.

D PRIMARY—The D module responsible for storing an X-Page's Data

D BACKUP—The D module responsible for storing a backup for an X-Page Data. The backup is stored in a non-volatile way (NVRAM or UPS protected).

ACRONYMS

LXA—Logical X-Page Address.

LB—Logical Block.

LBA—Logical Block Address.

AUS—Atomic Unit Size.

SL—Sub-LUN.

SLUS—Sub-LUN Unit Size.

MBE—Management Back End.

The examples described herein to a block-level storage system, offering basic and advanced storage functionality. The design may be based on a distributed architecture, where computational, Storage Area Networking (SAN), and storage elements are distributed over multiple physical Nodes, with all such Nodes being inter-connected over an internal network through a switch device. The distributed architecture enables the scaling of the system's capabilities in multiple aspects, including overall storage capacity, performance characteristics in bandwidth and I/O operations per second (IOPS), computational resources, internal and external networking bandwidth, and other. While being based on a distributed architecture, the system presents, externally, a unified storage system entity with scalable capabilities.

The system's architecture and internal algorithms implementing the basic and advanced storage functions are optimized for improved utilization of the capabilities of random-access memory/storage media, as opposed to contrast with mechanical-magnetic spinning disk storage media. The optimizations are implemented in the design itself, and may, for example, include the ability to break incoming writes into smaller blocks and distribute the operation over different Nodes. Such an adaptation is particularly suitable for random access memory/storage media but is less suitable in a spinning-disk environment, as it would degrade performance to extremely low levels. The adaptation includes the content/hash based mapping of data distributes the data over different D Nodes in general and within D Nodes over different SSD devices. Again, such a scheme is more suitable for random access memory/storage media than for a spinning-disk media because such spread of data blocks would result in very poor performance in the spinning disk case. That is to say, the described elements of the present architecture are designed to work well with random access media, and achieve benefits in performance, scalability, and functionality such as inline deduplication. Such random-access memory media can be based on any or a combination of flash memory, DRAM, phase change memory, or other memory technology, whether persistent or non-persistent, and is typically characterized by random seek/access times and random read/write speeds substantially higher than those exhibited by spinning disk media. The system's internal data block mapping, the algorithms implementing advanced storage functions, and the algorithms for protecting data stored in the system are designed to provide storage performance and advanced storage functionality at substantially higher performance, speed, and flexibility than those available with alternative storage systems.

Data mapping within the system is designed not only to improve performance, but also to improve the life span and reliability of the electronic memory media, in cases where the memory technology used has limitations on write/erase cycles, as is the case with flash memory. Lifetime maximization may be achieved by avoiding unnecessary write operations as will be explained in greater detail below. For the purpose of further performance optimization, life span maximization, and cost optimization, the system may employ more than a single type of memory technology, including a mix of more than one Flash technology (e.g., single level cell—SLC flash and multilevel cell—MLC flash), and a mix of Flash and DRAM technologies. The data mapping optimizes performance and life span by taking advantage of the different access speeds and different write/erase cycle limitations of the various memory technologies.

The core method for mapping blocks of data internally within the system is based on Content Addressing, and is implemented through a distributed Content Addressable Storage (CAS) algorithm.

This scheme maps blocks of data internally according to their content, resulting in mapping of identical block to the same unique internal location. The distributed CAS algorithm allows for scaling of the CAS domain as overall system capacity grows, effectively utilizing and balancing the available computational and storage elements in order to improve overall system performance at any scale and with any number of computational and storage elements.

The system supports advanced In-line block level deduplication, which may improve performance and save capacity.

Elements of the system's functionality are: Write (store) data block at a specified user address; Trim data block at a

specified user address; Read data block from a specified user address; and In-line block level deduplication.

The following features may be provided: (1) A distributed CAS based storage optimized for electronic random-access storage media; The optimization includes utilizing storage algorithms, mainly the content-based uniformly-distributed mapping of data, that inherently spread data in a random way across all storage devices. Such randomization of storage locations within the system while maintaining a very high level of performance is preferably achievable with storage media with a high random access speed; (2) A distributed storage architecture with separate control and data planes; Data mapping that maximizes write-endurance of storage media; System scalability; (3) System resiliency to fault and/or failure of any of its components; (4) Use of multi-technology media to maximize write-endurance of storage media; and (5) In-line deduplication in ultrahigh performance storage using electronic random-access storage media.

The examples described herein implement block storage in a distributed and scalable architecture, efficiently aggregating performance from a large number of ultra-fast storage media elements (SSDs or other), preferably with no performance bottlenecks, while providing in-line, highly granular block-level deduplication with no or little performance degradation.

One challenge is to avoid performance bottlenecks and allow performance scalability that is independent of user data access patterns.

The examples described herein may overcome the scalability challenge by providing data flow (Write, Read) that is distributed among an arbitrary and scalable number of physical and logical nodes. The distribution is implemented by (a) separating the control and data paths (the "C" and "D" modules), (b) maintaining optimal load balancing between all Data modules, based on the content of the blocks (through the CAS/hashing mechanisms), hence ensuring always balanced load sharing regardless of user access patterns, (c) maintaining optimal load balancing between all Control modules, based on the user address of the blocks at fine granularity, hence ensuring always balanced load sharing regardless of user access patterns, and (d) performing all internal data path operations using small granularity block size, hence detaching the incoming user access pattern from the internal access pattern, since the user pattern is generally larger than the block size.

A second challenge is to support inline, highly granular block level deduplication without degrading storage (read/write speed) performance. The result should be scalable in both capacity—which is deduplicated over the full capacity space—and performance.

The solution involves distributing computation-intensive tasks, such as calculating cryptographic hash values, among an arbitrary number of nodes. In addition, CAS metadata and its access may be distributed among an arbitrary number of nodes. Furthermore, data flow algorithms may partition read/write operations in an optimally-balanced way, over an arbitrary and scalable number of Nodes, while guaranteeing consistency and inline deduplication effect over the complete storage space.

In detaching the data from the incoming pattern, the R-Module breaks up any incoming block which is larger than the granularity size across sub-LUNs, sending the relevant parts to the appropriate C-Modules. Each C-module is predefined to handle a range or set of Sub-LUN logical addresses. The C-Module breaks up the block it receives for distribution to D-Modules, at a pre-determined granularity,

which is the granularity for which a Hash is now calculated. Hence the end result is that a request to write a certain block (for example of size 64 KB) ends up being broken up into for example 16 internal writes, each write comprising a 4 KB block.

The specific numbers for granularity can be set based on various design tradeoffs, and the specific number used herein of 4 KB is merely an example. The broken down blocks are then distributed to the D modules in accordance with the corresponding hash values.

A further challenge is to address flash-based SSD write/erase cycle limitations, in which the devices have a lifetime dependent on the number of write/erase cycles.

The solution may involve Inline deduplication to avoid writing in all cases of duplicate data blocks. Secondly, content (hash) based mapping to different data modules and SSDs results in optimal wear-leveling, ensuring equal spread of write operations to all data modules and SSDs independently of the user data/address access patterns.

In the following a system is considered from a functional point of view. As described above with respect to FIG. 1, the system **10** is architected around four main functional Modules designated R (for Router), C (for Control), D (for Data), and H (for Hash). Being modular and scalable, any specific system configuration includes at least one of R, C, D, and H, but may include a multiplicity of any or all of these Modules.

Reference is now made to FIG. 2, which is a functional block diagram of the system in which an H module **200** is connected to an R module **202**. The R module is connected to both Control **204** and data **206** modules. The data module is connected to any number of memory devices SSD **208**.

A function of the R Module **202** is to terminate SAN Read/Write commands and route them to appropriate C and D Modules for execution by these Modules. By doing so, the R Module can distribute workload over multiple C and D Modules, and at the same time create complete separation of the Control and Data planes, that is to say provide separate control and data paths.

A function of the C Module **204** is to control the execution of a Read/Write command, as well as other storage functions implemented by the system. It may maintain and manage key metadata elements.

A function of the D Module **206** is to perform the actual Read/Write operation by accessing the storage devices **208** (designated SSDs) attached to it. The D module **206** may maintain metadata related with the physical location of data blocks.

A function of the H Module is to calculate the Hash function value for a given block of data.

Reference is now made to FIG. 3, which illustrates nodes. The R, C, D, and H Modules may be implemented in software, and executed on a physical Node. A system includes at least one physical Node, and may include multiple Nodes. There are four possible Node configurations: Compute Node **300**, which includes control and hash modules, Compute+SAN Node **302** which includes a router as well as control and hash modules, Compute+Store Node **306**, which includes a data module in addition to compute and hash modules, and a Compute+SAN+Store Node **306**, which includes all four modules. A system includes a storage area networking or SAN function within at least one Node, and a Store function within at least one Node. The SAN function and the store function can be supported by the same physical Node or any combination of multiple Nodes.

In FIG. 3 each node type shows the functional Modules that execute, in at least one copy, within the Node, and

functional Modules that may optionally execute within this Node. Optional Modules are shown in dashed line.

All Nodes include a switch interface **308**, to allow interconnecting with a switch in a multi-Node system configuration. A Node that contains a SAN function includes at least one SAN Interface module **310** and at least one R Module. A Node that contains a Store function includes at least one SSD Driver Module **312** and at least one D Module. Hence, Compute+SAN and Compute+SAN+STORE Nodes contain a SAN Interface, to interface with the external SAN. The interface may typically use a SCSI-based protocol running on any of a number of interfaces including Fiber Channel, Ethernet, and others, through which Read/Write and other storage function commands are being sent to the system. Compute+Store and Compute+SAN+Store Nodes contain an SSD driver **312** to interface with SSDs **208** attached to that specific Node, where data is stored and accessed.

Reference is now made to FIG. 4, which shows a high level system block diagram. A system implementation includes one or more Nodes **400**, **402**. In all cases where a system contains more than two Nodes, all physical Nodes are interconnected by a switch **404** which may be based on any of a number of networking technologies including Ethernet, InfiniBand and so forth. In the specific case of a 2-Node system, the two Nodes can be interconnected directly without a need for a switch.

The interconnections between each Node and the Switch may include redundancy, so as to achieve high system availability with no single point of failure. In such a case, each Node may contain two or more Switch Interface modules **406**, and the Switch may contain two or more ports per physical Node.

As an example FIG. 5 illustrates a single Node system configuration, in which R, C and D modules, **500**, **502** and **504** respectively are together in a compute+SAN+Store node **506**. A switch interface **508** links to a switch. A SAN interface **510** provides an interface for storage area networking. An SSD driver **512** interfaces with the storage devices.

A four node system configuration is shown in FIG. 1 above. The configuration includes two compute and store nodes and two compute+SAN nodes.

A system that is built from multiple physical Nodes can inherently support a high availability construction, where there is no single point of failure. This means that any Node or sub-Node failure can be compensated for by redundant Nodes, having a complete copy of the system's meta-data, and a complete redundant copy of stored data (or parity information allowing recovery of stored data). The distributed and flexible architecture allows for seamless support of failure conditions by simply directing actions to alternate Nodes.

The R module is responsible for: routing SCSI I/O requests to the C modules, guarantee execution and return the result; and balancing the work load between the C modules for the requests it is routing.

An A→C table indicates which C module is responsible for each logical X-page address (LXA). Each C module is responsible for a list of Sub LUNs (SLs).

The R module receives requests for I/Os from the SAN INTERFACE, routes them to the designated C modules and returns the result to the SAN INTERFACE.

If an I/O operation spans across multiple SLs, and perhaps multiple C modules, then the R module has the responsibility of breaking the big I/O operation into multiple smaller independent operations according to the sub LUN unit size (SLUS). Since the atomic unit size (AUS) is never larger than the SLUS, as explained in greater detail below, each

such I/O is treated as an independent operation throughout the system. The results may then be aggregated before returning to the SAN INTERFACE.

The R module is responsible for maintaining an up-to-date A→C table coordinated with the MBE. The A→C table is expected to balance the range of all possible LXAs between the available C modules.

For write operations, the R module instructs the calculation of the hash digest for each X-Page by requesting such calculation from a Hash calculation module.

The C module is responsible for: receiving an I/O request from an R module on a certain SL, guaranteeing its atomic execution and returning the result; communicating with D modules to execute the I/O requests; monitoring the disk content of its SLs' logical space by associating each LXA with its hash digest; and balancing the work load between the D modules for the SLs it is maintaining.

An H→D table maps each range of hash digests to the corresponding D module responsible for this range.

An A→H table maps each LXA that belongs to the SLs C is responsible for, to the hash digest representing the X-Page Data that currently resides in this address.

The C module receives I/O requests from R modules, distributes the work to the D modules, aggregates the results and guarantees an atomic operation. The result is returned to the R module.

The C module maintains an up-to-date H→D table coordinated with the MBE. The table is expected to balance the range of all possible hash digests between the available D modules.

The C module maintains an A→H table in a persistent way. The C module may initiate **110** requests to D modules in order to save table pages to disk, and read them from disk. To avoid frequent disk operations, a Journal of the latest table operations may be maintained.

Data is balanced between the C modules based on the logical address, at the granularity of sub-LUNs.

The D module is responsible for: maintaining a set of LUNs which are attached locally and performing all I/O operations on these LUN; managing the physical layout of the attached LUNs; managing the mapping between X-Page Data hash digests and their physical location in a persistent way; managing deduplication of X-Page Data in a persistent way; and receiving disk I/O requests from C modules, perform them and returning a result.

The D module is also responsible for, for each write operation, backing up the X-Page Data in the designated D backup module and performing read-modify operations for writes that are smaller than X-Page size (This process also involves computing a hash digest for these X-Pages).

The D module is further responsible for maintaining an up-to-date H→(D, D_{backup}) table coordinated with the MBE. The H→(D, D_{backup}) table is expected to balance the range of all possible hash digests between the available D modules.

The D module does not communicate directly with R modules. The only interaction with R modules involves RDMA read/write operations of X-Page Data.

Balancing between the D modules is based on hashing of the content.

The D module makes use of a hash digest metadata table. The hash digest metadata table maps each in use hash digest, that represents actual X-Page Data, to its meta data information including its physical page on the storage media (SSD), its memory copy (if exists), a mapping to any backup memory copy and a reference count for the purpose of deduplication.

11

A further structure used is the $H \rightarrow (D, D_{\text{backup}})$ table. The $H \rightarrow (D, D_{\text{backup}})$ table maps each range of hash digests to the corresponding D module responsible for the range as well as the D_{backup} module responsible for the range.

The D modules allocate a physical page for each X-Page. The D modules also manage the memory for the physical storage. They allocate memory pages for read/write operations and perform background destaging from memory to storage media when necessary, for example, when running low on memory.

The D modules manage a separate nonvolatile memory pool (NVRAM or UPS protected) for X-Page Data backup purposes. The backup holds X-Pages that are held in memory of the D primary and have not yet been destaged. When re-balancing between D modules occurs (due to a D module failure for example), the D module may communicate with other D modules in order to create new backup copies or move a primary ownership as required.

The D modules allow deduplication per X-Page Data by maintaining a persistent reference count that guarantees only one copy per X-Page Data. The D modules manage the hash digest metadata table in a persistent way. The table is coordinated with the physical layout for physical pages allocation, with the memory pointer, memory backup pointer and deduplication reference count.

The D modules receive I/O requests from C modules, perform the requests while supporting deduplication and return the result. The D modules may perform RDMA read/write operations on memory that resides in other modules, such as R modules as mentioned above, as part of the I/O operation.

When a write operation smaller than the size of an X-Page is received, the D module may read the entire X-Page to memory and perform partial X-Page modification on that memory. In this case race conditions may occur, for example when two small writes to the same X-Page occur in parallel, and the D module may be required to compute the hash digest of the resulting X-Page. This is discussed in greater detail below.

The H-Module calculates the Hash function of a given block of data, effectively mapping an input value to a unique output value. The Hash function may be based on standards based hash functions such as SHA-1 and MD5, or based on a proprietary function. The hash function is selected to generate a uniformly distributed output over the range of potential input values.

The H modules usually share nodes with an R module but more generally, the H modules can reside in certain nodes, in all nodes, together with R modules, or together with C or D modules.

The following discussion provides high level I/O flows for read, write and trim.

Throughout these flows, unless noted otherwise, control commands are passed between modules using standard RPC messaging, while data "pull" operations may use RDMA read. Data push (as well as Journal) operations may use RDMA write.

The read flow of one X-Page may consist of one R module which receives the read request from the application, one C module in charge of the address requested and one D module which holds the X-Page to be read. Larger, or unaligned, requests may span several X-Pages and thus may involve several D modules. These requests may also span several SLs, in which case they may involve several C modules as well.

Reference is now made to FIG. 6 which illustrates the flow for a read operation for one aligned X-page. When the

12

R module receives a read request from an application the R module allocates a request ID for the operation; translates the LBA to LXA; allocates a buffer for the data to be read; consults the $A \rightarrow C$ component to determine which C module is in charge of this LXA; and sends the designated C module a read request which includes parameters that include a request ID; an LXA; and a pointer to the allocated buffer.

The C module, when receiving the request, consults the $A \rightarrow H$ component, from which it obtains a hash digest representing the X-Page to be read; consults the $H \rightarrow D$ component to determine which D module holds the X-Page in question; and sends this D module a read request which includes parameters that include a request ID (as received from the R module), the hash digest, a pointer to the buffer to read to, as received from the R module; and an identifier of the R module.

The D module, when receiving the request, reads the data of the requested X-Page from SSD and performs an RDMA write to the requesting R module, specifically to the pointer passed to it by the C module.

Finally the D module returns success or error to the requesting C module.

The C module in turn propagates success or error back to the requesting R module, which may then propagate it further to answer the application.

Reference is now made to FIG. 7, which illustrates the flow in the case that a read request arrives for a range of addresses spanning more than one X-Page but only one SL. In such a case the R module sends the designated C module a read command with the parameters that include a request ID, first LXA, size of the requested read in X-Pages-n, and n pointers to the allocated X-Page buffers.

The rest of the R module's treatment is identical to the aligned one X-Page scenario previously described herein.

The C module, when receiving the request divides the logical address space to LXAs. For each LXA the C module consults the $A \rightarrow H$ component to determine the corresponding hash digest; consults the $H \rightarrow D$ table to determine which D module is responsible for the current LXA; sends each D module a read command containing all the hashes that the respective D module is responsible for. The parameters of the read command include a request ID (as received from the R module); a list of respective hash-pointer pairs; and the identifier of the R module.

Each D module, when receiving the request, acts per hash-pointer pair in the same manner as described above for one X-Page. Aggregated success or error is then sent to the requesting C module.

The C module aggregates all the results given to it by the D modules and return success or error back to the requesting R module, which may then answer the application.

In the case that a read request spans multiple SLs, the R module splits the request and sends several C modules read requests. Each C module may receive one request per SL. The flow may continue as in the simpler case above, except that now the R module aggregates the responses before it answers the application.

Read requests smaller than 4 KB, as well as requests not aligned to 4 KB, may be dealt with at the R module level. For each such parcel of data, the R module may request to read the encompassing X-Page. Upon successful completion of the read command, the R module may crop the non-relevant sections and return only the requested data to the application.

The write flow of one X-Page may consist of one R module which receives the write request from the application, one C module in charge of the address requested and

13

three D modules: D_{target} which is in charge of the X-Page Data to be written (according to its appropriate hash digest), D_{old} which was in charge of the X-Page Data this address contained previously ("old" hash digest), and D_{backup} in charge of storing a backup copy of the X-Page Data to be written.

Reference is now made to FIG. 8, which is a simplified flow diagram illustrating the write procedure for a single aligned X page according to the examples described herein.

When an R module receives a write request from the application, the R module allocates a request ID for this operation; translates the LBA to an LXA; computes a hash digest on the data to be written; consults its A→C component to determine which C module is in charge of the current LXA; and sends the designated C module a write command with parameters that include a request ID; an LXA; a hash digest; and a pointer to the buffer containing the data to be written.

The C module, when receiving the request consults its H→D component to understand which D module is in charge of the X-Page to be written (D_{target}); and sends D_{target} a write request with parameters that include the request ID (as received from the R module); the hash digest (as received from the R module); the pointer to the data to write (as received from the R module); and the identifier of the R module.

The D module receiving the write command, D_{target} may first check if it already holds an X-Page corresponding to this hash. There are two options here:

First, D_{target} does not have the X-Page. In this case D_{target} fetches the data from the R module using RDMA read and stores it in its memory; consults the H→D component to determine which D module is in charge of storing a backup copy of this X-Page (D_{backup}); performs an RDMA write of the X-Page Data to the D_{backup} backup memory space; and returns success (or failure) to the C module.

Second, D_{target} has the X-Page. In this case D_{target} increases the reference count, returns success (or failure) to the C module.

The C module waits for a response from D_{target} . If a success is returned, the C module updates the A→H table to indicate that the LXA in question should point to the new hash and returns a response to the requesting R module.

If this is not a new entry in the A→H table, the C module asynchronously sends a decrease reference count command to D_{old} (the D module responsible for the hash digest of the previous X-Page Data). These commands may be aggregated at the C module and sent to the D modules in batches.

The R module may answer the application once it receives a response from the C module.

Reference is now made to FIG. 9, which is a flow diagram illustrating the process for writes to multiple full X-Pages.

In the case that the write request spans a range of addresses which include more than one X-Page but only one SL, the R module sends the designated C module a write command with parameters that include a request ID; a first LXA; a size of the requested write in LXAs-n; and H_{BIG} which is a unique identifier of the entire chunk of data to be written. H_{BIG} may be a computed hash digest and thus equal for two identical chunks of data.

Additional parameters sent with the write command are n pointers that point to the buffers which hold the data to be written.

The rest of the R module treatment is the same as for the aligned one X-Page scenario.

The C module, when receiving the request, consults its H→D component to understand which D module is in

14

charge of H_{BIG} (D_{target}) and generates a hash digest per pointer by replacing one byte of H_{BIG} with the offset of that pointer. It is noted that this byte must not collide with the bytes used by the H→D table distribution.

It may send D_{target} a write request with the parameters that include the request ID (as received from the R module); a list of respective hash-pointer pairs; and the Identifier of the R module.

The D module, when receiving the request, acts per hash-pointer pair in the same manner as described above for one X-Page. Aggregated success or error is then sent to the requesting C module.

The C module waits for a response from D_{target} . If the response indicates success, the C module updates its A→H table to indicate that the LXAs in question should point to the new hashes. Updating of entries in the A→H table may be done as an atomic operation, to ensure the write request is atomic. Note that all requests aligned to 4 KB (or another predefined block size) that fall within a SL may be atomic. The C module returns a response to the requesting R module. The C module adds the list of old hashes to the "decrease reference" batch if needed.

The R module answers the application once it receives a response from the C module.

In the case in which a write request spans multiple SLs, the R module splits the request and sends smaller write requests to several C modules. Each C module receives one request per SL (with a unique request ID). The flow continues as in the simpler case above, except that now the R module aggregates the responses before it answers the application.

Referring to FIG. 10, in one example, a D module 600 includes a hash digest metadata component (HMD) 602, which stores a hash digest metadata table 604, a user data cache component (UDC) 618 and a physical layout component (PLC) 632. The HMD 602 performs as a server to clients such as the UDC 618 and the PLC 632.

The hash digest metadata table 604 holds a reference count (ref-count) of the hash digest as part of deduplication implementation. The HMD 602 deletes a hash digest when its count reaches zero.

The UDC 618 component holds the data in memory and writes it asynchronously to the SSD using the PLC API. The UDC 618 also holds an in-memory read cache. The PLC 632 is responsible for managing the data on the SSD and for applying data protection algorithms such as RAID, for example. Both the UDC 618 and the PLC 632 use the HMD 602 to hold the location of the data per hash—in memory (UDC) or on SSD (PLC).

The HMD 602 performs mapping of hash digest to its metadata. The metadata is owned by the different clients as well as HMD 602 itself. HMD 602 also provides APIs 606 to generate new metadata entries, finds the new metadata entries and accesses their content. The HMD 602 also performs metadata backup for failure recovery. Some of the metadata that is kept per hash digest must be persistent. HMD 602 uses a journal. The journal is used to log metadata changes in remote memory and then a journal component (not shown) will aggregate these log entries and write them to the SSD in an amortized fashion. This approach is much faster than hardening every metadata change directly to SSD. The HMD 602 also hardens metadata to SSD to protect this data.

A hash digest takes up about 20 bytes. The expected number of hash digests in the hash digest metadata table 604 is limited by the amount of space available on the SSDs after excluding space reserved for metadata and data protection

(e.g., RAID). The challenge of managing a hash digest based table in a memory efficient way is not trivial, but having the HMD 602 manage the entire metadata allows for efficient processing. Thus, the client components (e.g., UDC 618 and PLC 632) are focused only with their own logic and algorithms. However, the client components have a well-defined interface with the HMD 602, which manages the hash digest based table 604 as well as the metadata backup. The HMD provide access to metadata of a specific hash digest given that hash or a handle that represents it.

There are different types of metadata that the HMD 602 is required to hold and there are challenges in keeping this metadata in a memory efficient way. For example, one challenge is to reduce as much as possible the amount of memory consumed per hash digest in order to reach a point where most of the metadata is held in memory and there is very little, if any, reading of any metadata from an SSD as part of an I/O flow.

Assuming that the hash digests values are distributed uniformly, the metadata entries are held in buckets where the first bytes of the digest are used as a bucket index. When determining how many buckets to use, there is a trade-off. The more buckets used, the less number of entries there is on average per bucket (better performance) but the more overall memory that is consumed.

In one example, the range of reasonable values is between using 24 bits (3 bytes) and using around 28 bits. As used herein 3 bytes are used as bucket index; however, other byte sizes may be used. Since the first 3 bytes of a hash digest are used as bucket index, there is no need to hold them inside each entry in that bucket, so that 3 bytes are saved in every entry.

A linked list of arrays or entries is held in a bucket. Each entry is marked by a flag if it is allocated or free to use (to avoid the memory of keeping a linked list of entries). When a new entry is allocated, the list is traversed to find a free spot—the list is traversed anyway to look for short hash collisions as further described herein. When an entry is free it is marked as free.

There is a tradeoff regarding what is the size of each array. If a small array is used more memory is wasted on the pointers between the arrays. If large arrays are used, more memory is wasted on “almost empty” arrays—where only one entry is in such an array, all the rest of that array is unused. Empty arrays are a larger waste, so that small arrays are used. There is no point to use an array smaller than a cache line (64 bytes), since cache lines will not be shared between buckets. Thus, a cache line as one array is chosen and a bucket will be a cache line chain. 4 bytes are taken out of every cache line as a pointer to the next cache line.

Every time a cache line that was linked to a chain is completely freed up from all its entries, it will be freed back to a large pool of cache lines, so that cache lines can move from bucket to bucket according to the distribution of entries between them.

When hardening the hash digest to SSD the entire 20 bytes of the hash digest are written. However, in memory less bytes can be held and any collisions are dealt with separately. This is called a short hash representation. In one example, the short hash is a short hash described in U.S. patent Ser. No. 14/037,626, filed Sep. 26, 2013 entitled “Generating a Short Hash Handle” assigned to the same assignee as the present patent application.

The total number of possible hash digests is limited by the SSD capacity. For 1 Tbyte, there are 256 M possible values. For 16 Tbytes, there are 4 Gbyte possible values.

6 bytes are used to hold a short hash representation. The first 3 bytes are used as bucket index. The remaining 3 bytes are used for managing short hash collisions. One bit will be used to mark if this short hash has a collision. This leaves 47 bits or 2^{47} possibilities. Thus, the collision probability for 16 Tbyte SSD capacity is $1:2^{15}$, meaning 1:32K. In the case of collisions where there is more than one entry with the same “short” hash, the first short hash that was added will be still held as a short hash. The second short hash will be put in a separate pool (array) of full hashes. This pool can hold up to 8 M hashes (2^{23}) since the suffix of 3 bytes (except for the collision bit) may be used as an index to this pool.

The size of the full hash pool is around 160 Mbyte. In cases where this pool is exhausted as well as cases where the full hash is required from a short hash, the full hash is read from the SSD and a read of the metadata is performed.

Referring to FIG. 11, for a metadata entry 700 there are persistent fields 714 and non-persistent metadata fields 730. The non-persistent fields 730 are owned by a client component, for example UDC 618, and are held by the client component itself. The HMD 602 metadata entry will have a cross reference against the corresponding metadata entry in the client component. Using the UDC 618 as an example (since UDC 618 has no persistent fields), the HMD 602 entry will hold a UDC (MD entry) handle that will allow the UDC 618 to get an metadata entry by hash digest. The UDC 618 will hold a handle to the HMD 602 entry to be able to modify the UDC handle when needed.

For persistent fields 714 it is more complex since the HMD 602 needs to backup the actual values of those fields, so it does require access to their values. In this case there are two options. The first option is to hold a client component handle (same as in the non-persistent case) and have an API obtain the actual values to harden when required. The second option is to hold the hardened fields directly in the metadata entry. In one example, the second option is chosen since it is simpler, and it avoids potential deadlocks. Deadlocks are cases where the HMD 602 is locking an entry A and calls a client component to get a metadata field value, and the client component tries to take lock B for this purpose. And in another thread, another client component locks lock B and then calls the HMD 602 which tries to lock entry A. Thus, the client components are waiting for each other. Since the HMD 602 is a server only and not a client requests come in one direction to avoid deadlocks.

Referring to FIG. 12, the persistent fields 714 include a short hash digest field 802, a physical layout offset field 808, a deduplication reference count field 814 and a flags field 822. The data saved in the persistent fields are saved in a persistent manner. The short hash digest field 802 is 3 bytes and is controlled by the HMD 602.

The physical layout offset field 808 is 4 bytes and 2 bits (34 bits) and is controlled by the PLC 632. The physical layout offset field 808 can be narrowed down to 34 bits assuming D module SSDs capacity is limited to 64 Tbytes and assuming the volume index is not coded inside the offset (“global” offset across different SSDs using some kind of mapping).

The deduplication reference count field 814 is 2.5 bytes. The maximal value here is dependent on the logical space of an application. Assuming, for example, 128 Tbyte will bring us to 2^{35} . In other examples up to 5 bytes could be safely used for this field. In order to achieve 2.5 bytes 20 bits are used where one bit marks the ones that overflow. If the bit is off, the remaining 19 bits hold the actual reference count. If the bit is on, the remaining 19 bits are an index to a separate array of larger ref-counts (5 bytes). The size of the

17

“large” ref-counts pool is a few Mbytes and we gain that in the common case we use only 2.5 bytes.

The flags field **822** is 0.5 bytes and is controlled by the HMD **602**. The flags field **822** includes two control flags on each entry. A first control flag marks if it is allocated or not, and the second control flag if it is extended or not (see below). These 2 bits can be combined with the 0.5 byte of the ref-count and can also hold the additional 2 bits of the physical layout offset

The number of entries that have valid non persistent fields **730** is significantly smaller than the total number of entries. The UDC **818** has such fields. The UDC size can be limited in 4 Gbytes, which is 1 M number of pages. This is a small enough number to hold a separate pool for “extended” entries with non-persistent fields.

The size of this pool can be 1 M entries and 3 bytes may be used as an index to each entry in the pool (array). In order not to add these 3 bytes to the metadata entry it is swapped with a persistent field and the field value is put in the pool entry. In one example, the physical layout offset may be chosen, since its 4 bytes. So, for entries that are currently extended, the extended entry index is put in the PL offset and the actual PL offset along with the non-persistent fields is in the extended entry.

Thus, a metadata entry contains 10 bytes. How much memory that is required may be estimated. For example, for a 1 Tbyte SSD gross capacity there are 256 M possible hash digests. If there are 16 M buckets, there is on average 16 entries per bucket. Since the metadata entries are put in cache lines, each 6 entries take 64 bytes. For 16 entries there are 3 cache lines which take 192 bytes. Therefore there is overall of 16 M*192=3 Gbyte of memory in the HMD **602** is consumed. For a 2 Tbyte SSD gross capacity 6 Gbyte of memory in the HMD **602** is consumed. For a 4 Tbyte SSD gross capacity 11 Gbyte of memory in the HMD **602** is consumed. For a 8 Tbyte SSD gross capacity 22 Gbyte of memory in the HMD **602** is consumed. For a 16 Tbyte SSD gross capacity 43 Gbyte of memory in the HMD **602** is consumed.

The system described herein may be implemented, at least in part, via a computer program product, (e.g., in a non-transitory machine-readable storage medium such as, for example, a non-transitory computer-readable medium), for execution by, or to control the operation of, data processing apparatus (e.g., a programmable processor, a computer, or multiple computers). Each such program may be implemented in a high level procedural or object-oriented programming language to communicate with a computer system. However, the programs may be implemented in assembly or machine language. The language may be a compiled or an interpreted language and it may be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment. A computer program may be deployed to be executed on one computer or on multiple computers at one site or distributed across multiple sites and interconnected by a communication network. A computer program may be stored on a non-transitory machine-readable medium that is readable by a general or special purpose programmable computer for configuring and operating the computer when the non-transitory machine-readable medium is read by the computer to perform the processes described herein. For example, the processes described herein may also be implemented as a non-transitory machine-readable storage medium, configured with a computer program, where upon execution, instructions in the computer program cause the computer to operate in accor-

18

dance with the processes. A non-transitory machine-readable medium may include but is not limited to a hard drive, compact disc, flash memory, non-volatile memory, volatile memory, magnetic diskette and so forth but does not include a transitory signal per se.

Elements of different embodiments described herein may be combined to form other embodiments not specifically set forth above. Other embodiments not specifically described herein are also within the scope of the following claims.

What is claimed is:

1. A method comprising:

providing a server that stores a hash digest metadata table, the hash digest metadata table mapping a hash digest representing X-Page data to its metadata information; configuring entries in the hash digest metadata table to include entries comprising persistent fields and non-persistent fields, the persistent fields comprising a short hash digest field, a physical layout offset field, a deduplication reference count field and a flags field;

configuring the server to be a server to clients comprising a user data cache component (UDC) and a physical layout component (PLC);

enabling the short hash digest field to be controlled by the server;

enabling the physical layout offset field to be controlled by the PLC;

enabling the deduplication reference count field to be controlled by the server; and

enabling the flags field to be controlled by the server.

2. The method of claim 1, further comprising configuring the short hash digest field to be 3 bytes in size.

3. The method of claim 1, further comprising configuring the physical layout offset field to be 4 bytes and 2 bits (34 bits).

4. The method of claim 1, further comprising configuring the deduplication reference count field is 2.5 bytes (20 bits).

5. The method of claim 1, further comprising configuring the flags field to be 0.5 bytes.

6. The method of claim 1, further comprising enabling the non-persistent field to be controlled by the UDC.

7. An article comprising:

a non-transitory computer-readable medium that stores computer-executable instructions, the instructions causing a machine to provide a server that stores a hash digest metadata table,

wherein the hash digest metadata table maps a hash digest representing X-Page data to its metadata information and the hash digest metadata table includes entries comprising persistent fields and non-persistent fields, the persistent fields comprising a short hash digest field, a physical layout offset field, a deduplication reference count field and a flags field;

wherein the server is configured to be a server to clients comprising a user data cache component (UDC) and a physical layout component (PLC), and

wherein the short hash digest field is controlled by the server, the physical layout offset field is controlled by the PLC, the deduplication reference count field is controlled by the server and the flags field is controlled by the server and the non-persistent field to be controlled by the UDC.

8. The article of claim 7, wherein the short hash digest field is 3 bytes in size.

9. The article of claim 7, wherein the physical layout offset field is 4 bytes and 2 bits (34 bits).

10. The article of claim 7, wherein the deduplication reference count field is 2.5 bytes (20 bits).

11. The article of claim 7, wherein configuring the flags field is 0.5 bytes.

12. An apparatus, comprising:

electronic hardware circuitry configured to provide a server that stores a hash digest metadata table, 5
wherein the hash digest metadata table maps a hash digest representing X-Page data to its metadata information, the hash digest metadata table comprising entries comprising persistent fields and non-persistent fields, the persistent fields comprising a short hash digest field, a 10
physical layout offset field, a deduplication reference count field and a flags field;

wherein the server is configured to be a server to clients comprising a user data cache component (UDC) and a physical layout component (PLC), and 15

wherein the short hash digest field is controlled by the server, the physical layout offset field is controlled by the PLC, the deduplication reference count field is controlled by the server and the flags field is controlled by the server and the non-persistent field to be controlled by the UDC. 20

13. The apparatus of claim 12 wherein the circuitry comprises at least one of a processor, a memory, a programmable logic device or a logic gate.

14. The apparatus of claim 12, wherein the short hash 25
digest field is 3 bytes in size, the physical layout offset field is 4 bytes and 2 bits (34 bits), the deduplication reference count field is 2.5 bytes (20 bits), the flags field is 0.5 bytes.

* * * * *